arXiv:1602.04379v1 [stat.AP] 13 Feb 2016

Adaptive Filter for Automatic Identification of Multiple Faults in a Noisy OTDR Profile

Jean Pierre von der Weid, Mario H. Souto, Joaquim D. Garcia, and Gustavo C. Amaral

February 16, 2016

Abstract

We present a novel methodology able to distinguish meaningful level shifts from typical signal fluctuations. A two-stage regularization filtering can accurately identify the location of the significant level-shifts with an efficient parameter-free algorithm. The developed methodology demands low computational effort and can easily be embedded in a dedicated processing unit. Our case studies compare the new methodology with current available ones and show that it is the most adequate technique for fast detection of multiple unknown level-shifts in a noisy OTDR profile.

1 Introduction

The central problem in fiber monitoring is the detection of small faults or losses most commonly performed by inspecting the trace of an Optical Time Domain Reflectometer (OTDR) [1]. These faults appear as small level shifts in a slowly varying backscattered optical power, eventually masked by the detector noise. Averaging over many OTDR shots is usually required to get access to the information needed. However, measurement time is of paramount importance in network monitoring, so that signal processing and filtering is a fundamental tool to improve time and sensitivity of the overall process. Moreover, in the case of wavelength multiplexed optical networks (WDM-PON) the problem is still worse because coherent backscattered power fluctuations (CRN) cannot be averaged out by summing up many OTDR shots [2].

The problem of identifying level shifts may be referred in the literature as *identification of structural breaks*, *step filtering*, *jump detection* or *regime switching* [3]. One possible approach to identifying such steps would be to decompose the signal into its level and additional components [4–6]. However, none of these methods can be directly applied to detecting level-shifts when the presence and position of the shift is unknown. Boyd and Kim [7] have introduced the ℓ_1 Trend Filter, a piece-wise linear filter capable of identifying location and magnitude of peaks and shifts. Recently, this technique has been successfully employed in automatically detecting faults in an optical fiber link [2].

In this paper, we analyze the potential of different techniques in identifying level shifts associated to fiber faults in a noisy OTDR profile and we make use of the Tunable OTDR

reported in [2] for the field tests. The main contributions of this paper are twofold. First, a three-steps extension to Boyd's ℓ_1 Trend Filter is proposed. The filter's accuracy on detecting the correct level-shift position is improved via a supplementary reweighted ℓ_1 minimization on the first two steps whereas the bias on the level shift estimate is corrected on the third. Since regularization leads to a shrunk estimator, a traditional least squares estimation of the selected level-shifts can produce unbiased magnitude estimators making the third step imperative. This novel three-steps filter is hereby referred to as the ℓ_1 *Adaptive Filter*.

The second contribution consists of an efficient algorithm for the first two steps of the ℓ_1 Adaptive Filter. Each of those steps is a convex optimization problem and there are multiple off-the-shelf techniques for solving it. Nevertheless, the special structure of the problem induces a tailored algorithm. Inspired by the Covariance Update Coordinate Descent [8], we developed a faster filtering algorithm designed for this particular problem.

This paper is organized as follows. We present the test bench optical fiber link used throughout the article and its profile acquired by the Tunable OTDR in Section 2. Section 3 starts by describing the level-shift filter. It is shown how the filtered signal can be obtained by solving a special minimization problem that leads to the Least Absolute Shrinkage and Selection Operator (LASSO) [9], which is the starting point of our tailored algorithm. Section 4 introduces an adaptive and improved version of the filter based on a reweighted ℓ-1 minimization and proposes an efficient algorithm for solving the adaptive filter. The comparison between several signal processing techniques and our own is exposed in Section 5.1 where the focus is on the ability of fast and accurate detection of level shifts in the test bench profile. Finally, in Section 6, conclusions are drawn over the effectiveness of our novel methodology and some future researches are suggested.

2 Optical fiber fault detection

The Optical Time Domains Reflectomettry (OTDR) technique consists of sending a light pulse into an optical fiber and measuring the Rayleigh backscattered light [1]. The OTDR trace is the backscattered power as a function of the distance: a descending line in logarithmic scale with angular coefficient equal to the fiber's attenuation coefficient [10]. Any event which causes optical power loss is interpreted accordingly so fiber faults and defective fiber splices can be identified by abrupt level shifts in the signal profile. Due to its mathematical characteristics, the signal can, therefore, be decomposed into piecewise linear functions with slopes that correspond to the attenuation coefficient of the respective fiber stretch. Recently, a Tunable OTDR has been proposed in which Boyd's ℓ_1 Trend Filter is used to perform such decomposition and the positions and magnitudes o the faults are automatically identified.

The achieved spatial resolution reported in [2] is 5.72 meters with a spatial indexation of the data series of 1 meter due to the hardware maximum clock speed limitation. We employed this technique to measure the profile of a test bench optical fiber link with several fault events specifically designed as to mimic the difficulty in monitoring an unknown fiber link. The testbench fiber has approximately 12 kilometers and was tailored to include interesting features such as the presence of small faults next to big ones, several faults spaced by no more than 100 meters, reflective and non-reflective fault events, and fibers

with different attenuation coefficients. The OTDR profile is depicted in Fig. 1.

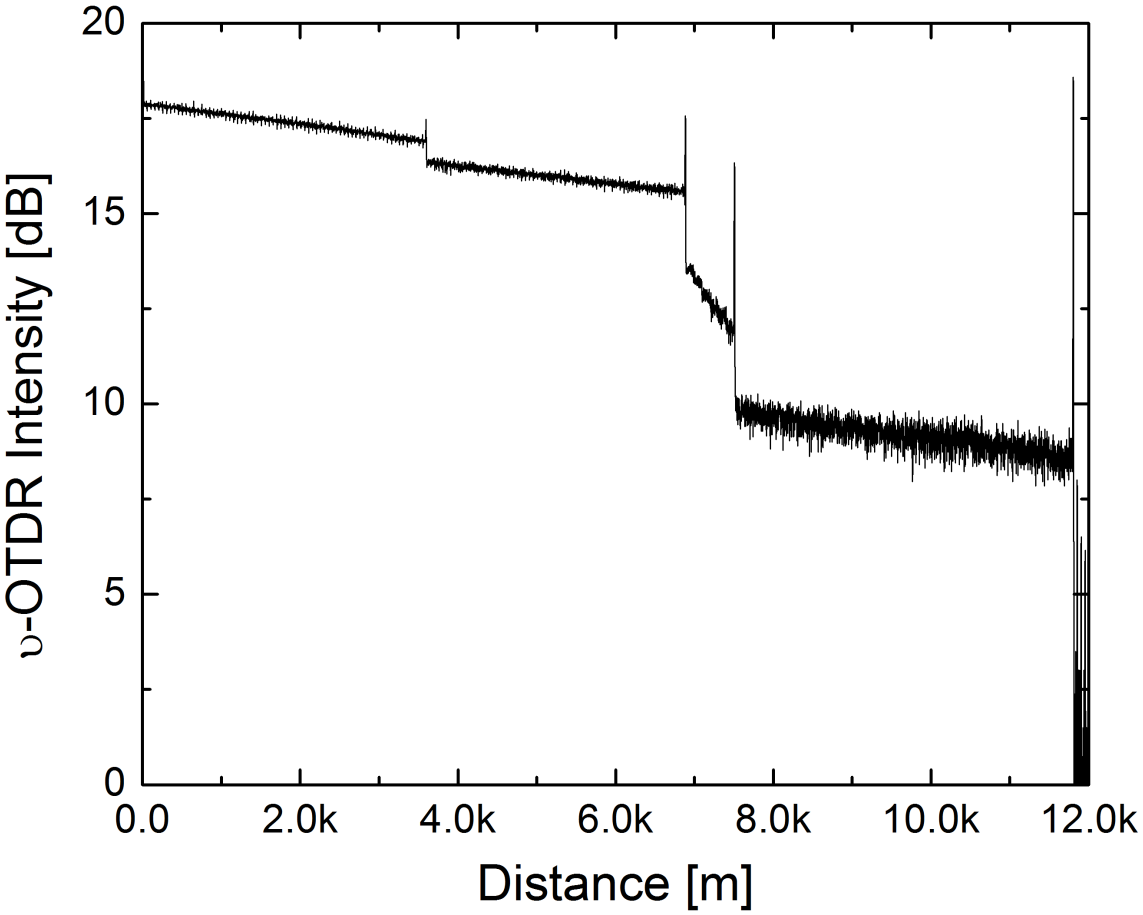

Figure 1: Raw ν-OTDR profile of the testbench 12 kilometers fiber. Several faults spaced by no more than 100 meters were induced between 6 and 8 kilometers. The first fiber stretch up to $\sim$ 3.5 kilometers has a different attenuation coefficient which is translated by a linear region with different slope.

Two main sources of noise are present in Tunable-OTDR measures: the Poisson Noise associated with the counting process; and the Coherent Rayleigh Noise (CRN) associated with the random interference from backscattered wave fronts. Since the first grows with the square root of the number of detections, the SNR due to Poisson Noise can be set at arbitrary levels by extending the monitoring period. This approach, however, would lead to monitoring periods above the expected from an OTDR technique. The CRN, on the other hand, is an intrinsic effect which cannot be fully eliminated since it is dependent on the probe pulse's linewidth. As commented in [11], wavelength sweep techniques can be employed to minimize its contribution but not to eliminate it. In this context, a filtering method which deals with such noise sources mathematically without compromising the technique's performance is highly desirable.

3 Filtering via the ℓ_1 norm

As introduced in [7], the main idea of the ℓ_1 Trend Filter is to decompose the signal into a piecewise linear curve. To do so, the filter aims to minimize the sum of squared residuals simultaneously penalizing the first difference in the filtered signal to obtain a sparse solution; in this case, sparsely distributed level-shifts.

The most intuitive idea is to directly minimize the number of level-shifts. Unfortunately, this is a combinatorial task and belongs to the class of NP-hard problems [12]. With the development of theoretical results [13] in the past twenty years, the idea of using the ℓ_1 norm in place of the counting norm has gained much attention and acceptance [14] [15] . Consider a vector $y \in \mathbb{R}^N$ containing the signal observations and let N be the total number of observations. We can define the filtered signal as the combination of a piecewise-linear component and occasional level-shifts. The filtering process can be achieved by solving the following multi-objective optimization problem:

$$\min_{(a,\{\theta_k\}_{k=1}^N)} \frac{1}{2}\sum_{k=1}^{N}(y_k - \theta_k - ak)^2 + \lambda \sum_{k=2}^{N} |\theta_k - \theta_{k-1}| \tag{1}$$

In this sense, the signal y is decomposed into the sum of a monotonic linear function with the slope $a \in \mathbb{R}$ and the level $\theta \in \mathbb{R}^N$. The penalty term $\lambda \sum_{k=2}^{N} |\theta_k - \theta_{k-1}|$, penalizes shifts on the level component. Thus, the regularization encourages the level to be as constant as possible. This filter can be seen as a particular form of the fused lasso [16].

The tuning parameter λ is a non-negative real number that controls the frequency and the amount of components used by the filter to describe the signal of interest. By changing the value of λ it is possible to manage the trade-off between the total sum of residuals and the number of components of the filtered signal. The optimization problem (1) can be reduced to the problem of finding a sparse solution for a particular minimization problem. Let us consider the matrix $X \in \mathbb{R}^{N\times N}$ which includes all the possible candidates for a fault position in its columns, i.e., the j^{th} candidate is a *step* function of the form $u_{-1}(k-j)$ which is described, in matrix form, as a vector of ones up to the j^{th} position and zeros in the remaining positions. The simplest form of the matrix X which considers a single slope is presented in Eq. 2.

$$X = \begin{bmatrix} 1 & 0 & 0 & \cdots & 0 & 0 \\ 2 & 1 & 0 & \cdots & 0 & 0 \\ 3 & 1 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & & \ddots & \vdots & \vdots \\ N-1 & 1 & 1 & \cdots & 1 & 0 \\ N & 1 & 1 & \cdots & 1 & 1 \end{bmatrix} \tag{2}$$

When written as $\mu = X\beta$, the filtered signal μ is interpreted as the linear combination of selected level shifts which best approximate the original signal y. The vector of unknowns $\beta \in \mathbb{R}^N$ corresponds to the component selector: whenever the i^{th} position of β is different from zero, a level shift component is introduced at position i in the linear combination to form μ. Consequently, the signal y can be stated as $y = X\beta + \varepsilon$, where

ε is the approximation error. With the aforementioned re-parametrization, it becomes clear that the filter (1) is equivalent to a regularized least squares of the form:

$$\min_{(\beta_i)} \tfrac{1}{2}\|X\beta - y\|_2^2 + \lambda \sum_{i=2}^{N} |\beta_i| \tag{3}$$

The expression in Eq. (3) is a method of major interest in several fields of study. The most renowned is the LASSO [9], which is used to perform model fitting and variable selection in the high-dimensional statistics framework. Another well-known method is the *Basis Pursuit Denoising* [17], its equivalent in the signal processing literature, broadly used in sparse signal recovery.

4 ℓ_1 Adaptive Filter

The main idea of the ℓ_1 Adaptive Filter is to firstly elect potential candidates by finding the vector β_j^I which minimizes Eq. (3). Afterwards, each of those elected variables are differently penalized by some parameter that is proportional to $\frac{1}{|\beta_j^I|^\gamma}$ for some $\gamma > 0$. Finally, the selected variables, or components, are given by

$$\min_{\beta} \tfrac{1}{2}\|X\beta - y\|_2^2 + \lambda \sum_{2}^{N} \frac{1}{|\beta_i^I|^\gamma} |\beta_i| \tag{4}$$

Since the selected components are comprised within the ones elected by Eq. (3), fewer components are set to be different from zero by the ℓ_1 Adaptive Filter.

In order to select the best estimators, a two dimensional grid, usually built on logarithmic scale, is created for the pair $\{\rho, \gamma\}$. The best pair needs to be chosen by some metric evaluated at every point of the two dimensional grid. According to Zou [18], the *Bayesian Information Criterion* (BIC) [19] can consistently select the values of the regularizer that will most likely recover the correct set of variables [20].

Zou proposed the *Adaptive* LASSO [21] as an improvement on the ability of identifying significant variables for the problem described by Eq. 4. The greatest benefit is the technique's consistency in variable selection since the underlying true model should be identified when the number of observations grows. Nevertheless, the LASSO produces biased estimators due to the shrinking process behind the technique [22][23]. As a result, the filter's capability to perform the signal recovery is diminished. More precisely, the magnitude of every level-shift will be smaller than expected. This last issue can be easily solved by calculating the Ordinary Least Square (OLS) estimation after the adaptative filter has chosen the proper components to describe the signal. Since just a few components will be nonzero, obtaining the OLS estimator will be straightforward and computationally cheap.

A series of algorithms are known to solve problems of the form presented in Eqs. (3) or (4). Since the two problems are very similar apart from the weighting factor in the latter, they can be solved by the same technique. It is also important to highlight that both optimization problems are convex and therefore any local minimum must be a global minimum [24]. We propose a tailored version of the *Covariance Updates Coordinate Descent* [8][25] for solving the optimization problem.

The Coordinate Descent algorithm is very popular in the field of high-dimensional statistics due to its excellent computational times. The core of the method is to optimize a multivariate function by sequentially minimizing the objective function with respect to a single coordinate direction instead of the multivariate direction given by the gradient [26]. Since the objective function is composed of a differentiable and convex component plus a non-differentiable but separable component, optimality is guaranteed [27]. The optimization process works, in every iteration, as to minimize the objective function at a coordinate without changing the values of the remaining ones. The algorithm will perform the core cycle of Algorithm 1 until no coordinate has changed its value which means convergence was reached.

Algorithm 1 Coordinate Descent

Initialize active set $\mathcal{A} = \varnothing$
for all $\lambda_k \in \{\lambda_1, \cdots, \lambda_{max}\}$ **do**
  while not converge **do**
    for all $j \in \{1, .., N-1\}$ **do**
      $\hat{\beta}_j^{ols} = \frac{1}{N}\left[\langle X_j, y\rangle - \sum_{l \in \mathcal{A}} \langle X_j, X_l\rangle \tilde{\beta}_l\right] + \tilde{\beta}_j$
      if $\alpha_{k,i,j} \geq |\hat{\beta}_j^{ols}|$ **then**
        $\tilde{\beta}_j = 0$
        $\mathcal{A} = \mathcal{A} \backslash \{j\}$
      else
        $\tilde{\beta}_j = sign(\hat{\beta}_j^{ols})(|\hat{\beta}_j^{ols}| - \alpha_{k,i,j})$
        if $j \notin \mathcal{A}$ **then**
          $\mathcal{A} = \mathcal{A} \cup \{j\}$
        end if
      end if
    end for
  end while
  $\hat{\beta}_j^{CD}(\lambda_k) = \tilde{\beta}_j \quad \forall \ j = 0, .., N-1$
end for

The Coordinate Descent method presented above can take different forms, namely *Type 1* and *Type 2*, depending on the expression assumed by $\alpha_{k,i,j}$. *Type 1* is characterized by $\alpha_{k,i,j} = \lambda_k \sigma_j$ whereas, for *Type 2*, $\alpha_{k,i,j} = \frac{\lambda_k \sigma_j}{|\beta_j^I| \gamma_i}$. This is the same as assuming $\beta_j^I = 1$ for the *Type 1* case. As discussed extensively in [8], the computational effort of each cycle depends on the entry of a new component X_j in the model which is accompanied by the computation of the inner products $\langle X_j, X_l\rangle$ ($O(N)$ operations) and the update of the current coefficients in the model ($O(p)$ operations) where p is the number of components in the model. Hence, a model that can be described by m components will take $O(Nm)$ operations.

Taking advantage of the special structure of the matrix X, the inner products necessary for the run of Algorithm 1 can be analytically determined as a function of the values of X. We dedicate Appendix I to the mathematical development of such analytic expressions and other features as the calculation of the inner products between the columns of the

matrix X and the vector y. This procedure, also necessary for the run of Algorithm 1, can be executed in a simplified and more efficient manner whenever the implementation has enough memory space available. The storage of the computed inner products in a Gram Matrix, on the other hand, does not to limit the effort associated to the procedure therefore the algorithm can be ran with equal efficiency in a device with limited available memory such as a Micro-Controller or an FPGA.

Algorithm 1 is used to obtain the filter's estimate for a given tuning parameter λ where the potential values for λ are within the range $[0, \lambda_{MAX}]$. The fact that λ_{MAX} can be calculated in beforehand renders the algorithm completely parameter-free. Let σ be the vector containing the standard deviation of each column of matrix X. Then,

$$\lambda_{MAX} = \left\{ \lambda = \left| \sum_{i=1}^{N} X_{ij} y_i \right| : \lambda \geq \left| \sum_{i=1}^{N} \frac{\sigma_k}{\sigma_j} X_{ik} y_i \right| \forall\, k \neq j \right\} \tag{5}$$

When $\lambda = 0$, the solution will be the same as the OLS estimate with a great amount of level-shifts. On the other hand, when $\lambda = \lambda_{MAX}$, there will be no components to describe the signal. These are the two extreme cases of the LASSO grid. It is important to note that the calculations of λ_{MAX} should be performed before the normalization discussed in Appendix I. The final solution algorithm for the ℓ_1 Adaptive Filter is presented in Algorithm 2 for which Algorithm 1 is the core.

Algorithm 2 ℓ_1 Adaptive Filter

Initialize $\tilde{\beta}_j = 0 \quad \forall\ j = 0, .., N-1$

Run Algorithm 1 Type 1.
Obtain: $\hat{\beta}_j^{CD}(\lambda_k) \quad \forall\ j = 0, .., N-1$ *and* $k = 1, .., m$.
Define: $\hat{\beta}^I$ as the best $\beta^{CD}(\lambda_k)$

for all $\gamma_i \in \{\gamma_1, \cdots, \gamma_q\}$ **do**
  Run Algorithm 1 Type 2.
  Obtain: $\hat{\beta}_j^{ADA}(\lambda_k, \gamma_i) = \hat{\beta}_j^{CD}(\lambda_k)$
  $\forall\ j = 0, .., N-1$ *and* $k = 1, .., m$.
end for
Define: $\hat{\beta}$ as the best $\hat{\beta}^{ADA}(\lambda_k, \gamma_i)$
Compute: OLS estimator for $\hat{\beta}^{ADA}(\lambda_k, \gamma_i) \neq 0$

For simplicity, the algorithm is initialized with $\hat{\beta}_j^{CD}(\lambda_k) = 0$, $\forall \lambda_k \in \{\lambda_1, \cdots, \lambda_m\}$. This enhances the efficiency of the algorithm since it saves many comparisons and memory accesses. The interpretation for this event is that the j^{th} value of β_j^I was chosen to be irrelevant during initialization.

5 Results and Discussion

In Fig. 2, we show the filtered OTDR profile when the ℓ_1 Adaptive Filter is employed. The algorithm is capable of accurately selecting the induced faults, even those as small as 0.1 dB. Another feature of the algorithm is that faults close to each other are unequivocally detected as detailed by the inset of Fig. 2. This result is important to attest the sensibility of the method even in the presence of consecutive small faults next to big faults.

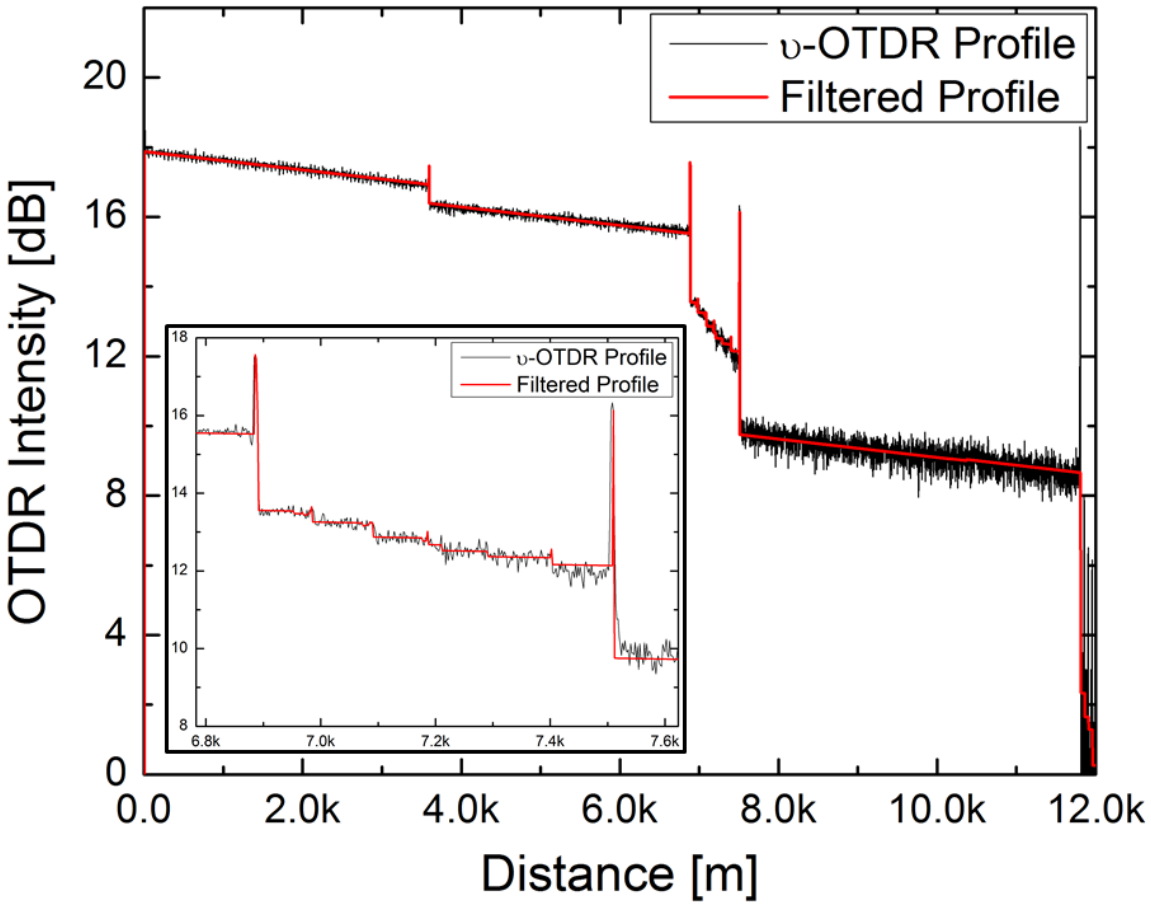

Figure 2: Raw ν-OTDR Profile (black) and ℓ_1 Adaptive Filter result (red). The inset details five consecutive faults ranging from 0.1 to 0.3 dB within a 600 meters stretch which were accurately detected by the algorithm.

5.1 Comparison between different Methodologies

From the result of Fig. 2, it becomes clear that the ℓ_1 Adaptive Filter represents a potential candidate for optical fiber fault detection applications. However, whether or not it is the best candidate depends on the performance of others techniques aimed at the same purpose. Therefore, we study the ℓ_1 Adaptive Filter capacity on detecting fiber faults against recent mathematical approaches such as the Potts Functional [28], and classical ones such as the ART [29] and the ℓ_1-style minimization [7, 30] which was the technique presented in [2] in conjunction with the Tunable OTDR for automatic fault

detection. A total of six different methodologies were compared: the L1 Potts functional; the L2 Potts Functional; the reweighted L1 minimization proposed in [30]; the ART using BIC as a pruning criterion; the LASSO; and the proposed ℓ_1 Adaptive Filter.

The L1 Potts, the L2 Potts, the reweighted L1 and the LASSO depend on a tuning parameter so a grid of size 100 was created and the best model was chosen via the BIC. The ℓ_1 Adaptive Filter uses the LASSO estimator as an initial condition therefore it is subject to the LASSO grid. A simpler version of the problem, with no linear trend, was used as input for both the L1 and L2 Potts algorithms since they do not account for slopes. Comparisons between the methodologies were drawn in the light of four main parameters: elapsed time; number of unsuccessful detections (within the resolution of the technique); and number of spurious detections. These were elected given the ultimate goal of a fiber monitoring method: fast and accurate detection of a fault.

The ν-OTDR acquisition period is intimately tied to the error associated to the number of count events at each position due to the Poisson Noise as commented in Section 2. A reasonable SNR is therefore necessary so the competing algorithms can truly be tested in its prowess to identify fiber characteristics in an OTDR profile. The testbench fiber presented in Figs. 1 and 2 demands a minimum of 100 seconds acquisition period in order for the last portion of the fiber link to be well above the noise floor in the OTDR profile which correspond to approximately $8 \cdot 10^5$ launched probing pulses. The rate at which the monitoring technique launches the probe pulses depend on the fiber's length due to a restriction of 1 pulse at a time traversing the fiber [2]. In the presented case, the rate is 8 kHz for the ~ 12 kilometer-long testbench fiber.

The acquired data sets from 100 seconds up to 300 seconds ($8 \cdot 10^5$ to $2.4 \cdot 10^6$ launched probe pulses respectively) – with 1 second spacing between each sample – were fed into the listed algorithms so their performance could be assessed. The graphs depicted in Figs. 3 and 4 trace the algorithms elapsed time and the number of spurious detections versus number of probing pulses launched into the testbench fiber. Due to the difficulty in graphical visualization, the number of missed faults is not displayed in a separate graph, but rather along with the results of the spurious detections in Fig. 4.

A few comments should be included before the results are discussed. The Reweighted L1 is unable to detect all the consecutive faults between 6 and 8 kilometers. The Potts L2, despite being able to detect some of these faults, fails to detect all of them. These results are independent on the acquisition time. Both the LASSO and the ART algorithms detect all the induced faults with minimum acquisition time. The Potts L1 and ℓ_1 Adaptive Filter get more sensitive as the acquisition time rises but the first takes as long as ~ 250 seconds to detect all of the faults whereas the second reaches this mark with ~ 200 seconds.

Analysis of Fig. 3 indicates that the ℓ_1 Adaptive Filter shows comparable timing characteristics to the fastest methods with no more than 10 seconds difference. The increase in elapsed time as the acquisition time rises, even though counter intuitive at first, can be attributed to the convergence criterion of the *Coordinate Descent*: as the faults become more distinctive, the selection between two candidates separate by few meters gets harder due to numerical issues. There are methods which deal with such numerical impairments and can be employed to accelerate the algorithm's convergence []. The Potts L1 and L2 algorithms exhibit higher than one minute elapsed times and its curves were, therefore, subtracted from the graph to ease the visualization.

Fig. 4 conveys a good deal of information regarding the accuracy of each method.

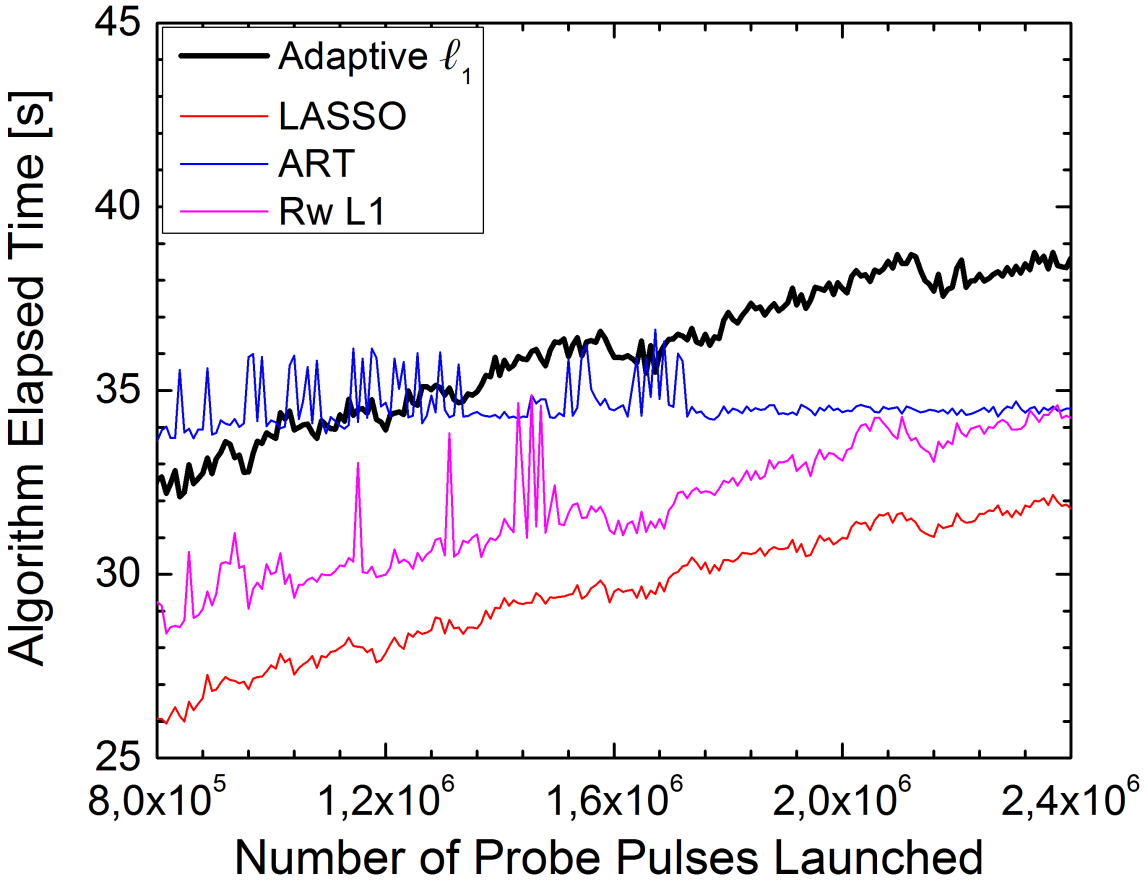

Figure 3: Algorithm elapsed time versus the number of probing pulses launched into the testbench fiber. Probe pulse rate is 8 kHz. The L1 and L2 Potts Functional yield higher than one minute processing time so their results were subtracted from the graph to ease the visualization.

Those methodologies which achieved the lowest number of missed faults, i.e., required the least acquisition time to identify all the faults in the link, namely the LASSO and ART algorithms, are also the ones with the highest number of spurious detections as depicted by the *red* and *blue* lines of Fig. 4. Conversely, the Reweighted L1 algorithm shows the smallest number of spurious detections but is unable to detect the five consecutive fault events which are detailed in the inset of Fig. 2.

Data acquisition, as commented in Section 2, is performed in 1 meter bins even though the spatial resolution of the Tunable OTDR technique is 5.72 meters. For that reason, a fault event is scarcely characterized by a single step within a 1 meter slot but rather represented by a cluster of steps inside a 6 meter range. In order to simplify our comparison method, all fault positions identified by the algorithms which do not match the value indicated by the standard OTDR profile serving as reference, are considered spurious. Therefore, the apparent high number of spurious detections of the ℓ_1 Adaptive Filter is actually an expected result. Considering an average of 3 spurious detections per cluster (smaller faults tend to be represented by less than a 6 meter clusters), we find that the ℓ_1 Adaptive Filter hardly identifies a step outside a cluster and figures as the most accurate

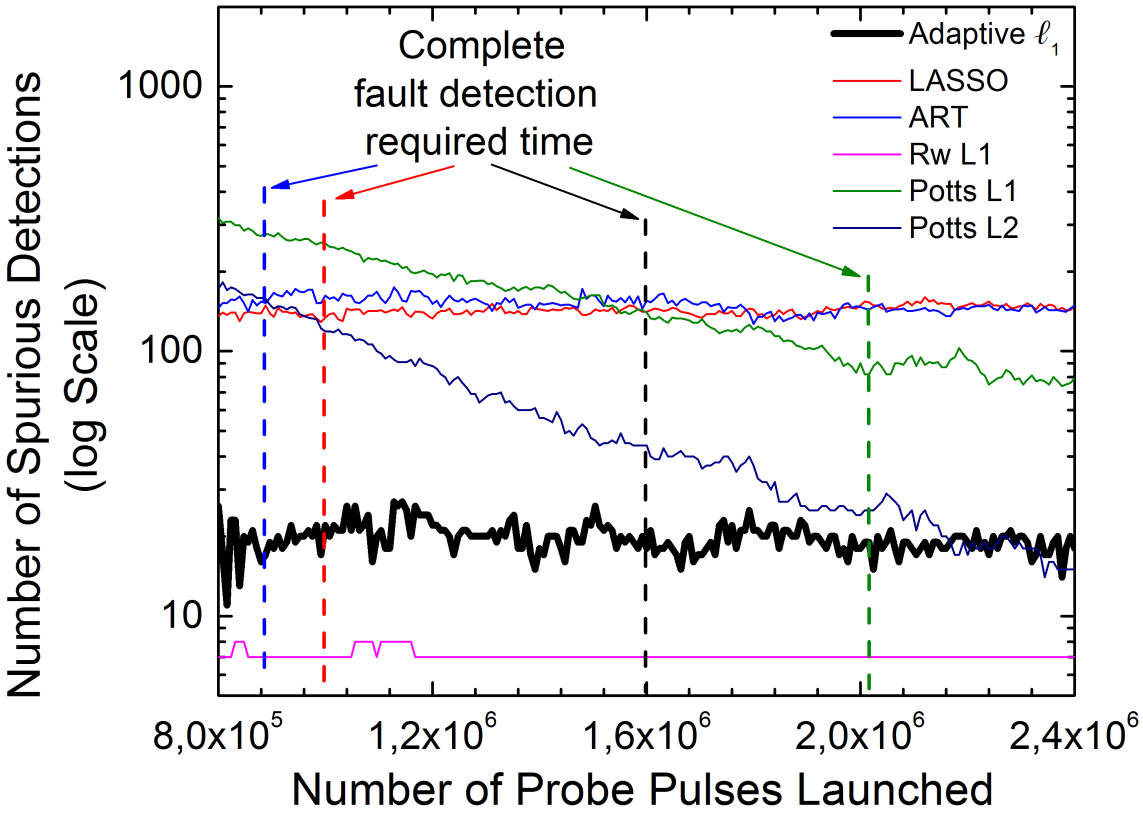

Figure 4: Number of spurious detections versus the number of probing pulses launched into the testbench fiber. Probe pulse rate is 8 kHz. This feature represents the accuracy of the filter in detecting the faults of the testbench fiber. The traced lines correspond to the minimum time required by each technique for detecting all faults present in the testbench link. The Reweighted L1 and Potts L2 algorithms were not able to identify all faults within 300 seconds of acquisition time.

technique among those presented. The association of level shift clusters with a fault has already been dealt with in [2].

For the presented testbench, the ℓ_1 Adaptive Filter was able to include all faults in the optical link from 200 seconds acquisition time with the benefit of a minimum number of spurious detections. The algorithm's precision, accuracy, competitive timing (total monitoring time including acquisition and processing steps is ~ 245 seconds), and readily applicability in the architecture presented in [2] sets the ℓ_1 Adaptive Filter as a promising candidate for noisy OTDR filtering and fault detection method. The ℓ_1 Adaptive Filter presents the best compromise between the number of spurious detections and the ability to identify the faults present in the link with reasonable processing time.

6 Conclusion

We have shown that the proposed ℓ_1 Adaptive Filter outperforms well-established methods found in the literature that tackle the problem of detecting level shifts in a

data stream. The main contribution of our study is a novel successful tool for detecting multiple level-shifts at unknown times in a signal of interest with direct applications in the filtering and automatic identification of faults in a noisy OTDR profile.

Additionally, we proposed an algorithm based on the Coordinate Descent which offers enhanced selection accuracy and competitive timing. Inspection of a testbench fiber optical link profile with different method shows that the ℓ_1 Adaptive Filter presents the best compromise between accurate and fast detection of fault events. The low computational effort demanded by the algorithm allows it to be embedded in dedicated processing units such as an FPGA or Micro-Controller.

Among the many possible future research directions we mention two: i) Inclusion of a spike trend filter, along with the step trend filter, in order to discern between a reflective and a non-reflective fiber fault and also to eliminate pitfalls at the final portion of the fiber which contains several spurious detections.; ii) Exploration of the parallel computing capabilities of the Coordinate Descent algorithm to increase the time response of the filter.

Appendix I

It is important to observe that the matrix X, as defined in Eq. 2, is not fit to be fed into Algorithm 1 since its columns ought to be statistically normalized in order to have comparable candidates of explanatory variables. In order to do so, we have to subtract each column of X by its mean and divide each of them by their own standard deviation. This normalization has an analytic form as follows.

We label the lines and columns of the $N \times N$ matrix X with the numbers 0 to $N-1$. Both mean and standard deviation of the i^{th} column are respectively given by:

$$\mu_i = \begin{cases} \frac{N+1}{2} & \text{if i = 0,} \\ \\ \frac{N-i}{N} & \text{otherwise.} \end{cases}$$

$$\sigma_i = \begin{cases} \sqrt{\frac{1}{6}(N+1)(2N+1) - \mu_0^2} & \text{if i = 0,} \\ \\ \sqrt{\frac{1}{N}(i\mu_i^2 + (N-i)(1-\mu_i)^2)} & \text{otherwise.} \end{cases}$$

Then, we define $U_i = \frac{-\mu_i}{\sigma_i}$ and $L_i = \frac{1-\mu_i}{\sigma_i}$, for $i \neq 0$ with the following interpretation: U_i and L_i are the values of the positions occupied by zeroes and ones on the column i of X, respectively, after the above mentioned normalization. Since U is an upper diagonal and L a lower diagonal matrix, both can be stored simultaneously in a single $N \times N$ matrix. Finally, the terms that comprise the Gram Matrix GM, i.e., the matrix containing the inner products between every pair of columns of X, are given by:

$$GM_{00} = N - 1$$

$$GM_{0j} = \frac{1}{\sigma_0}\left[jU_j\left[\frac{j+1}{2} - \mu_0\right] + (N-j)L_j\left[\frac{j+1+N}{2} - \mu_0\right]\right]$$

$$GM_{j0} = GM_{0j}$$

$$GM_{ij} = min(i,j)U_iU_j + (max(i,j) - min(i,j))L_iU_j +$$

$$(N - max(i,j))L_iL_j$$

Therefore, depending on the size of the problem and the physical memory available for the algorithm, there are two possibilities: calculate the values with the above formulas as they become necessary, or pre-compute the Gram Matrix and store the results. These closed formulas are much faster than computing all the inner products in standard way, so there is gain in time for both approaches.

The run time of Algorithm 1 can also be optimized if one pre-computes the values $\langle X_j, y\rangle$, the inner products between each column of the normalized X matrix and the normalized y vector. The inner products $\langle X_j, y\rangle$ can be more efficiently computed by calculating $\langle X_0, y\rangle$ in standard manner and then using Algorithm 1 for $\langle X_i, y\rangle$.

Algorithm 3 Fast y Inner-Products

Initialize $sumY = 0$
for i **from** 1 **to** $N - 1$ **do**
  $sumY = sumY + y_i$
  $\langle X_j, y\rangle = sumY(U_i - L_i)$
end for

Acknowledgment

The authors wish to thank brazilian agencies CPNq, CAPES and FAPERJ for financial support.

References

[1] M. K. Barnoski, M. D. Rourke, S. M. Jensen, and R. T. Melville, "Optical Time Domain Reflectometer," *Applied Optics*, vol. 16, 1977.

[2] G. Amaral, J. Dias Garcia, L. Herrera, G. Temporao, P. Urban, and J. von der Weid, "Automatic fault detection in wdm-pon with tunable photon counting otdr," *Lightwave Technology, Journal of*, vol. 33, no. 24, pp. 5025–5031, 2015.

[3] R. S. Tsay, "Outliers, level shifts, and variance changes in time series," *Journal of forecasting*, vol. 7, no. 1, pp. 1–20, 1988.

[4] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Journal of Fluids Engineering*, vol. 82, no. 1, pp. 35–45, 1960.

[5] J. Durbin and S. J. Koopman, *Time series analysis by state space methods*. Oxford University Press, 2012, no. 38.

[6] R. J. Hodrick and E. C. Prescott, "Postwar us business cycles: an empirical investigation," *Journal of Money, credit, and Banking*, pp. 1–16, 1997.

[7] S.-J. Kim, K. Koh, S. Boyd, and D. Gorinevsky, "ℓ_1 trend filtering," *Siam Review*, vol. 51, no. 2, pp. 339–360, 2009.

[8] J. Friedman, T. Hastie, and R. Tibshirani, "Regularization paths for generalized linear models via coordinate descent," *Journal of statistical software*, vol. 33, no. 1, p. 1, 2010.

[9] R. Tibshirani, "Regression shrinkage and selection via the lasso," *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 267–288, 1996.

[10] G. P. Agrawal, *Fiber-Optic Communication Systems*, 2nd ed. John Wiley & Sons, 2002.

[11] D. Villafani Caballero, J. von der Weid, and P. Urban, "Tuneable otdr measurements for wdm-pon monitoring," in *Microwave Optoelectronics Conference (IMOC), 2013 SBMO/IEEE MTT-S International*, 2013, pp. 1–5.

[12] B. K. Natarajan, "Sparse approximate solutions to linear systems," *SIAM journal on computing*, vol. 24, no. 2, pp. 227–234, 1995.

[13] D. L. Donoho and P. B. Stark, "Uncertainty principles and signal recovery," *SIAM Journal on Applied Mathematics*, vol. 49, no. 3, pp. 906–931, 1989.

[14] J. B. Rosen, H. Park, J. Glick, and L. Zhang, "Accurate solution to overdetermined linear equations with errors using l1 norm minimization," *Computational Optimization and Applications*, vol. 17, no. 2–3, pp. 329–341, 2000.

[15] W. Oberhofer, "The consistency of nonlinear regression minimizing the l1-norm," *The Annals of Statistics*, vol. 10, no. 1, pp. 316–319, 1982.

[16] R. Tibshirani, M. Saunders, S. Rosset, J. Zhu, and K. Knight, "Sparsity and smoothness via the fused lasso," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 67, no. 1, pp. 91–108, 2005.

[17] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM journal on scientific computing*, vol. 20, no. 1, pp. 33–61, 1998.

[18] H. Zou, T. Hastie, R. Tibshirani *et al.*, "On the "degrees of freedom" of the lasso," *The Annals of Statistics*, vol. 35, no. 5, pp. 2173–2192, 2007.

[19] G. Schwarz *et al.*, "Estimating the dimension of a model," *The annals of statistics*, vol. 6, no. 2, pp. 461–464, 1978.

[20] Y. Zhang, R. Li, and C.-L. Tsai, "Regularization parameter selections via generalized information criterion," *Journal of the American Statistical Association*, vol. 105, no. 489, pp. 312–323, 2010.

[21] H. Zou, "The adaptive lasso and its oracle properties," *Journal of the American statistical association*, vol. 101, no. 476, pp. 1418–1429, 2006.

[22] P. Bühlmann and S. Van De Geer, *Statistics for high-dimensional data: methods, theory and applications*. Springer, 2011.

[23] T. Hastie, R. Tibshirani, J. Friedman, T. Hastie, J. Friedman, and R. Tibshirani, *The elements of statistical learning*. Springer, 2009, vol. 2, no. 1.

[24] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2009.

[25] J. Friedman, T. Hastie, H. Höfling, R. Tibshirani *et al.*, "Pathwise coordinate optimization," *The Annals of Applied Statistics*, vol. 1, no. 2, pp. 302–332, 2007.

[26] P. Tseng and S. Yun, "A coordinate gradient descent method for nonsmooth separable minimization," *Mathematical Programming*, vol. 117, no. 1-2, pp. 387–423, 2009.

[27] P. Tseng *et al.*, *Coordinate ascent for maximizing nondifferentiable concave functions*. Massachusetts Institute of Technology, Laboratory for Information and Decision Systems, 1988.

[28] M. Storath, A. Weinmann, and L. Demaret, "Jump-sparse and sparse recovery using potts functionals," 2014.

[29] W. S. Rea, M. Reale, C. Cappelli, and J. A. Brown, "Identification of changes in mean with regression trees: an application to market research," *Econometric Reviews*, vol. 29, no. 5-6, pp. 754–777, 2010.

[30] E. J. Candes, M. B. Wakin, and S. P. Boyd, "Enhancing sparsity by reweighted $\ell 1$ minimization," *Journal of Fourier analysis and applications*, vol. 14, no. 5-6, pp. 877–905, 2008.